



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

AF 2FW
2195

Applicant: Mingjiu Sun et al.

Title: FAULT-TOLERANT SYSTEM AND METHODS WITH TRUSTED MESSAGE ACKNOWLEDGMENT (As Amended)

Docket No.: 884.439US1
Filed: June 7, 2001
Examiner: Kenneth Tang

Serial No.: 09/876,645
Due Date: October 11, 2005
Group Art Unit: 2195

MS Appeal Brief - Patents

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

We are transmitting herewith the following attached items (as indicated with an "X"):

2

- ☒ Return postcard.
- ☒ APPEAL BRIEF UNDER 37 CFR 41.37, including Appendices (26 pgs.).
- ☒ Authorization to charge Deposit Account No. 19-0743 in the amount of \$500 for Appeal Brief Fee.

If not provided in a separate paper filed herewith, Please consider this a **PETITION FOR EXTENSION OF TIME** for sufficient number of months to enter these papers and please charge any additional fees or credit overpayment to Deposit Account No. 19-0743.

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
Customer Number 21186

By: Ann M. McCrackin
Atty: Ann M. McCrackin
Reg. No. 42,858

CERTIFICATE UNDER 37 CFR 1.8: The undersigned hereby certifies that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail, in an envelope addressed to: MS Appeal Brief - Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on this 11th day of October, 2005.

CAROLYN HULSEY
Name

Carolyn Hulsey
Signature

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
(GENERAL)



PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants:	Mingqiu Sun et al.	Examiner:	Kenneth Tang
Serial No.:	09/876,645	Group Art Unit:	2195
Filed:	June 7, 2001	Docket No.:	884.439US1
Title:	FAULT-TOLERANT SYSTEM AND METHODS WITH TRUSTED MESSAGE ACKNOWLEDGMENT		
Assignee:	Intel Corporation	Customer No.:	21186

APPEAL BRIEF UNDER 37 CFR § 41.37

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is presented in support of the Notice of Appeal to the Board of Patent Appeals and Interferences, mailed on August 11, 2005, from the Final Rejection of claims 1-36 of the above-identified Application, as set forth in the Final Office Action mailed on March 11, 2005, and the Advisory Action mailed on June 27, 2005.

The Commissioner of Patents and Trademarks is hereby authorized to charge Deposit Account No. 19-0743 in the amount of \$500.00 which represents the requisite fee set forth in 37 C.F.R. § 41.2(b)(2). The Appellant respectfully requests consideration and reversal of the Examiner's rejections of the pending claims.

10/14/2005 EFLORES 00000106 190743 09876645

01 FC:1402 500.00 DA



APPEAL BRIEF UNDER 37 C.F.R. § 41.37

TABLE OF CONTENTS

	<u>Page</u>
<u>1. REAL PARTY IN INTEREST</u>	2
<u>2. RELATED APPEALS AND INTERFERENCES</u>	3
<u>3. STATUS OF THE CLAIMS</u>	4
<u>4. STATUS OF AMENDMENTS</u>	5
<u>5. SUMMARY OF CLAIMED SUBJECT MATTER</u>	6
<u>6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL</u>	9
<u>7. ARGUMENT</u>	10
<u>8. SUMMARY</u>	18
<u>CLAIMS APPENDIX</u>	19
<u>EVIDENCE APPENDIX</u>	25
<u>RELATED PROCEEDINGS APPENDIX</u>	26

1. REAL PARTY IN INTEREST

The real party in interest of the above-captioned patent application is the assignee, INTEL CORPORATION.

2. RELATED APPEALS AND INTERFERENCES

There are no other appeals, interferences, or judicial proceedings known to Appellant that will have a bearing on the Board's decision in the present appeal.

3. STATUS OF THE CLAIMS

The present Application was filed on June 7, 2001 with claims 1-31. After a Non-Final Office Action and a Final Office Action (hereinafter “the Final Office Action”) were received, with timely responses filed thereto, an Advisory Action was mailed on June 27, 2005.

At this time, claims 1-36 are currently pending in the Application. Claims 1-36 stand rejected, and their rejection is appealed herein.

4. STATUS OF AMENDMENTS

In a Preliminary Amendment dated October 18, 2001, claims 22 and 23 were amended, and new claims 32-36 were added. No further amendments have been made to the claims, and no additional new claims have been added.

After the Final Office Action, Appellant filed a Response, without any amendments, dated June 13, 2005, which was indicated as “entered” in the Examiner’s Advisory Action dated June 27, 2005. However, the Examiner indicated in the Advisory Action that Appellant’s request for reconsideration did not place the application in condition for allowance.

5. SUMMARY OF CLAIMED SUBJECT MATTER

This summary is presented in compliance with the requirements of Title 37 C.F.R. § 41.37(c)(1)(v), mandating a “concise explanation of the subject matter defined in each of the independent claims involved in the appeal ...”. Nothing contained in this summary is intended to change the specific language of the claims described, nor is the language of this summary to be construed so as to limit the scope of the claims in any way.

Some embodiments of the invention relate to fault-tolerant methods performed by a data processing system (claim 1), by a computer network (claims 9 and 22), by a computer (claim 17), by a computer-readable medium (claim 27), and by an article comprising a machine-accessible medium (claim 32).

A computer network 1 (Application FIG. 1) may comprise a plurality of nodes 2-7. Node 2 may be a fault-tolerant system 2 (Application FIG. 2, described on page 4, lines 15-23). Fault-tolerant system 2 may include at least two computers 11 and 12, each of which may also be referred to as a “workflow engine”. Each workflow engine includes software 27 (Application FIG. 2, described on page 5, lines 19-24) that may include middleware software and workflow recovery software.

In an embodiment, a data processing system comprises a number (M) of client processes 51-53 (Application FIG. 3) and a number (N) of workflows 71-73. The client processes 51-53 may reside on one or more of nodes 2-7 of network 1 (page 5, lines 27-29). The workflows 71-73 may reside on one or more of nodes 2-7 of network 1 (page 6, lines 18-20). A load manager 60 (Application FIG. 3, described beginning on page 5, line 30, through page 7, line 6) comprises a distributed queue 62 and a confirmation manager 64. Utilizing the distributed queue 62, load manager 60 balances the processing load, including requests for workflow execution originating from client processes, among the available workflow engines (page 6, lines 26-29).

In an embodiment (as described on page 7, lines 22-30), a fault-tolerant method includes a client process (e.g. client process A) requesting a workflow be performed by sending a request message RM_A to load manager 60. Load manager 60 may eventually assign the client-originating request to one of workflows 71-73 (e.g. workflow 72) by

sending a workflow assignment message WA_A to workflow 72. Workflow 72 may perform one or more tasks to carry out the desired client-requested functions (page 8, lines 1-2).

When, and only when, workflow 72 concludes its final task T_{FT} , workflow 72 sends a finished message FIN_2 to confirmation manager 64 of load manager 60 (page 8, lines 2-6), which sends an acknowledgement message ACK_A to client process A notifying it that workflow 72 has been performed (page 8, lines 7-9). Acknowledgement message ACK_A may be referred to as a “delayed acknowledgement message”, because it is only sent to client process A upon the successful completion of final task T_{FT} by workflow 72 (page 8, lines 11-13). Such delayed acknowledgement message can be trusted by the client process A, because the client process A can assume that the requested function has been successfully completed (page 8, lines 13-15).

If a software or hardware failure occurs in workflow 72, or if for any reason messages WA_A or FIN_2 are not delivered, load manager 60 will eventually realize this fact, and load manager 60 will subsequently assign the client-originating workflow request to a different workflow, such as workflow 71 or workflow 73 (page 8, lines 18-23).

Some embodiments of the invention are recited in independent claims 1, 9, 17, 22, 27, and 32, which will now be discussed in two different groups below.

Independent claims 1 and 9 recite operations that read on the operations of the flow diagram illustrated in Application FIG. 5 (described beginning page 9, line 29 through page 10, line 20). In 102, distributed queuing of workflows is provided among a plurality of workflow engines. In 104, a determination is made whether a workflow whose execution has been requested by a client has been completed by a first workflow engine. If so, the method goes to 106; otherwise, it goes to 108. In 106, an explicit and delayed acknowledgement is sent to the execution-requesting client, and the method ends. In 108, the workflow is assigned to another workflow engine.

Independent claims 17, 22, 27, and 32 recite operations that read on the operations of the flow diagram illustrated in Application FIG. 6 (described beginning page 10, line 20 through page 11, line 6). In 202, a workflow execution is requested on behalf of a

client. In 204, the workflow execution is assigned to a first workflow engine. In 206, a determination is made whether the workflow has been completed by the first workflow engine. If so, the method goes to 208; else, it goes to 210. In 208, an explicit and delayed acknowledgement is sent to the client, and the method ends. In 210, the workflow is assigned to another workflow engine.

6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

6.1 Claims 1-36 stand rejected under 35 U.S.C. §112, second paragraph, as being incomplete for “omitting essential structural cooperative relationships of elements, such omission amounting to a gap between necessary structural connections”.

6.2 Claims 1-36 stand rejected under 35 USC §103(a) as being unpatentable over Campbell et al. (U.S. 2001/0024497 A1) in view of Nakamura et al. (U.S. 2001/0027477 A1).

7. ARGUMENT

7.1 35 U.S.C. §112, SECOND PARAGRAPH REJECTION

7.1.1 The Applicable Law

35 U.S.C. §112, second paragraph, states:

“The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.”

In reviewing a claim for compliance with 35 U.S.C. §112, second paragraph, the examiner must consider the claim as a whole to determine whether the claim apprises one of ordinary skill in the art of its scope and, therefore, serves the notice function required by 35 U.S.C. §112, second paragraph, by providing clear warning to others as to what constitutes infringement of the patent. M.P.E.P. § 2173.02, citing *Solomon v. Kimberly-Clark Corp.* 216 F.3d 1372, 1379, 55 USPQ2d 1279, 1283 (Fed. Cir. 2000) and *In re Larsen*, No. 01-1092 (Fed. Cir. May 9, 2001) (unpublished). In the latter case, the court observed that the totality of all the limitations of the claim and their interaction with each other must be considered to ascertain the inventor’s contribution to the art.

7.1.2 The Examiner’s Rejection and Appellant’s Arguments

Claims 1, 9, 17, 22, 27, and 32 were rejected under 35 U.S.C. §112, second paragraph, as being incomplete for “omitting essential structural cooperative relationships of elements, such omission amounting to a gap between the necessary structural connections”.

The Examiner made five points in his rejection, each of which will be discussed in turn below:

1) The Examiner asserted that the “omitted structural cooperative relationships” are “one or more execution-requesting clients” and “a plurality of workflow engines”.

In response, Appellant asserts that regarding claim 1, for example, a clear cooperative relationship between the “one or more execution-requesting clients” and “a plurality of workflow engines” is set forth in the first operation or “step”. However, as an aid to the Board in more fully understanding claim 1, Appellant will briefly elaborate on the elements of claim 1 and their cooperative relationship.

In claim 1, a method is to be performed by a data processing system. The data processing system comprises a plurality of workflow engines. Thus, there are two or more workflow engines. The method provides distributed queuing of workflows among the plurality of workflow engines. Thus, there is more than one workflow. Execution of workflows is requested by one or more execution-requesting clients. Thus, there is at least one execution-requesting client in the data processing system.

Appellant asserts that the Examiner’s statement that the limitations “one or more execution-requesting clients” and “a plurality of workflow engines” are “omitted structural cooperative relationships” is indefinite and unsupported, because these limitations are already both clearly recited in claim 1, thus providing ample warning to one of ordinary skill in the art regarding the scope of claim 1. Thus, Appellant asserts that claim 1 fully complies with 35 U.S.C. §112, second paragraph.

2) The Examiner asserted that it is not made explicitly clear in the claim language whether the first client is linked only to the first workflow engine, and the second client linked only to the second workflow engine, and so on.

In response, Appellant asserts that regarding claim 1, for example, the answer to this issue is that the first client is not linked only to the first workflow engine, nor is the second client linked only to the second workflow engine, etc. In the first operation or “step” of claim 1, there is no one-to-one fixed linkage recited or intended to be recited regarding the clients and the workflow engines. Further, the second operation or “step” of claim 1 explicitly states that “if a workflow is completed by a first workflow engine for an execution-requesting client, sending an explicit and delayed acknowledgement to the execution-requesting client, else assigning the workflow to a second workflow engine”, so again no one-to-one fixed linkage exists or is intended to exist between the

clients and the workflow engines. Thus, Appellant's reply is that the first client is not linked only to the first workflow engine, and that a second client is not linked only to the second workflow engine. If, for a particular client, completion of the workflow execution fails on a first workflow engine, then execution is transferred to a second workflow engine. If execution fails on the second workflow engine, execution is transferred to a third workflow engine, and so forth, up to the limit of workflow engines in the data processing system.

The relationship between execution-requesting clients and workflow engines in claim 1 would be clear to one of ordinary skill in the art. Thus, Appellant asserts that claim 1 fully complies with 35 U.S.C. §112, second paragraph.

3) The Examiner asserted that, in other words, it is unclear whether or not there is a workflow engine for every client.

In response, Appellant asserts that regarding claim 1, for example, the answer to this issue is that the data processing system contains "M" clients (see FIG. 3) and "N" workflow engines" (see FIG. 2). Thus, the number of clients can range between 1 and M, and the number of workflow engines can range between 2 and N. The answer to the Examiner's question is that every client has at least one workflow engine available to its use, and it can have a second workflow engine assigned to it, if necessary. However, every client need not have a unique workflow engine dedicated solely to its use.

Again, as mentioned earlier, the relationship between execution-requesting clients and workflow engines in claim 1 would be clear to one of ordinary skill in the art. Thus, Appellant asserts that claim 1 fully complies with 35 U.S.C. §112, second paragraph.

4) The Examiner asserted that it is unclear whether or not the first client can communicate with the second workflow engine.

In response, Appellant respectfully asserts that regarding claim 1, for example, the answer to this issue is that it is not relevant to the subject matter recited in claim 1. Nothing in claim 1 describes any communication by a first client, either with a first

workflow engine or with a second workflow engine. Nor is any such communication required for the operation of the inventive subject matter of claim 1.

Appellant asserts that the relationship between a first client and the second workflow engine in claim 1 would be clear to one of ordinary skill in the art. Thus, Appellant asserts that claim 1 fully complies with 35 U.S.C. §112, second paragraph.

5) The Examiner asserted that it is unclear whether or not there is the same number of workflow engines as clients.

In response, Appellant asserts that regarding claim 1, for example, the answer to this issue is found above in the answer to question 3). That is, the number of clients can range between 1 and M, and the number of workflow engines can range between 2 and N. Thus the answer to this issue is there need not be the same number of workflow engines as clients.

Again, Appellant asserts that the numerical relationship between clients and workflow engines in claim 1 would be clear to one of ordinary skill in the art. Thus, Appellant asserts that claim 1 fully complies with 35 U.S.C. §112, second paragraph.

Regarding independent claims 9, 17, 22, 27, and 32, the same responses apply as were provided above regarding independent claim 1.

For the above reasons, Appellant respectfully requests that the rejection of claims 1, 9, 17, 22, 27, and 32 under 35 U.S.C. §112, second paragraph, be withdrawn.

7.2 35 U.S.C. §103 REJECTION

7.2.1 *The Applicable Law*

The Examiner has the burden under 35 U.S.C. §103 to establish a *prima facie* case of obviousness. *In re Fine*, 837 F.2d 1071, 1074, 5 U.S.P.Q.2d (BNA) 1596, 1598 (Fed. Cir. 1988). The M.P.E.P. contains explicit direction to the Examiner that agrees with the *In re Fine* court:

In order for the Examiner to establish a *prima facie* case of obviousness, three base criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure. *M.P.E.P.* § 2142 (citing *In re Vaeck*, 947 F.2d 488, 20 U.S.P.Q.2d (BNA) 1438 (Fed. Cir. 1991)).

The requirement of a suggestion or motivation to combine references in a *prima facie* case of obviousness is emphasized in the Federal Circuit opinion, *In re Sang Su Lee*, 277 F.3d 1338; 61 U.S.P.Q.2D 1430 (Fed. Cir. 2002), which indicates that the motivation must be supported by evidence in the record.

The test for obviousness under §103 must take into consideration the invention as a whole; that is, one must consider the particular problem solved by the combination of elements that define the invention. *Interconnect Planning Corp. v. Feil*, 774 F.2d 1132, 1143, 227 U.S.P.Q. 543, 551 (Fed. Cir. 1985). References must be considered in their entirety, including parts that teach away from the claims. See MPEP § 2141.02.

The Examiner must avoid hindsight. *M.P.E.P.* § 2143.01 (citing *In re Gordon*, 733 F.2d 900, 221 U.S.P.Q. 1125 (Fed. Cir. 1984)). That is, the Examiner cannot use the Appellant's structure as a "template" and simply select elements from the references to reconstruct the claimed invention. *In re Gorman*, 933 F.2d 982, 987, 18 U.S.P.Q.2d (BNA) 1885, 1888 (Fed. Cir. 1991). If the proposed modification renders the prior art invention being modified unsatisfactory for its intended purpose, then there is no

suggestion or motivation to make the proposed modification. *M.P.E.P.* § 2143.01 (citing *In re Gordon*, 733 F.2d 900, 221 U.S.P.Q. 1125 (Fed. Cir. 1984)).

7.2.2 *The References*

Campbell discloses (see Abstract) a system and method for servicing multi-media customer communications to geographically distributed agents from multiple call center sites via the telephone network and a global data communications network. Appellant could find no disclosure in Campbell concerning sending an explicit and delayed acknowledgement to an execution-requesting client if a workflow is completed by a first workflow engine, else assigning the workflow to a second workflow engine. It is noted that the Examiner concedes this in his Final Office Action.

Nakamura discloses (see Abstract) a workflow-controlling system comprising a single server 110 (FIG. 1) and a plurality of client terminals 130, 140, and 150. Nakamura states that “a flow control section that is a central component of the present invention exists on server 110, and documents are managed by this section” (Para. 72). Thus, Nakamura discloses only a single workflow engine. Nakamura’s “terminals” are merely client apparatus (Para. 72), not workflow engines.

7.2.3 *Discussion of the §103 Rejections*

Claims 1-36 were rejected under 35 U.S.C. §103(a) as being unpatentable over Campbell et al. (U.S. 2001/0024497 A1) in view of Nakamura et al. (U.S. 2001/0027477 A1).

First, the Appellant does not admit that Campbell or Nakamura are prior art, and reserves the right to swear behind these references in the future.

Second, since a *prima facie* case of obviousness has not been established, the Appellant respectfully traverses this rejection.

In this Appeal, Appellant is not asserting (1) lack of motivation to combine the references, or (2) lack of reasonable expectation of success if the references are combined. Appellant is solely asserting (3) no combination of the references discloses all the limitations set forth in the claims.

The References Do Not Teach All Claim Limitations.

As mentioned earlier, the Examiner concedes that Campbell fails to explicitly teach sending an explicit and delayed acknowledgement to an execution-requesting client if a workflow is completed by a first workflow engine, else assigning the workflow to a second workflow engine.

However, the Examiner stated that Nakamura teaches determining whether a first workflow is completed by a first client/terminal and, if complete, a notification is sent; if not completed, then a different notification is sent so that the second terminal/client can be assigned the workflow, citing Paras. 11 and 12.

However, Appellant respectfully asserts that the Examiner is not construing Nakamura's disclosure correctly in Paras. 11 and 12.

Setting aside for the moment the fact that Nakamura discloses only one workflow engine, Appellant asserts that Nakamura's basic operation is fundamentally different from that recited in claim 1. Nakamura is not concerned with fault tolerance but with work flow. Nakamura is not concerned with completion of an entire workflow by a single workflow engine and notifying the requesting client of this fact.

The "update request" operation described in Nakamura's Paras. 11 and 12 is described in detail in Para. 98 regarding FIGS. 9 and 10. It will be noted that Nakamura uses the term "completed", regarding a workflow, to mean "finished" or "concluded" with respect to all clients, not just one client. Thus Para. 98 explains that if an "update request" is not completed, then one or more additional terminal/clients/participants (i.e. participants #2 and #3, FIG. 10) are notified that they may provide input. This a very different operation from that described in Appellant's claim 1, in which the entire workflow (see paragraph immediately below) is assigned to a second workflow engine if the workflow is not completed by a first workflow engine for a client.

Appellants concede that claim 1 doesn't explicitly recite "entire workflow". However, FIG. 4 of the Application defines a "workflow" to include a starting task and a final task, as well as one or more intermediate tasks. In Nakamura, the work isn't necessarily completed by one participant, so a partial workflow may be worked on by subsequent participants. Appellant's claim 1 states that if a workflow (i.e. a starting task,

one or more optional intermediate tasks, and a final task) is completed by a first workflow engine for an execution-requesting client, sending an explicit and delayed acknowledgement to the execution-requesting client, else assigning the workflow (i.e. a starting task, one or more optional intermediate tasks, and a final task) to a second workflow engine.

Further, Appellant could find no disclosure in Nakamura concerning sending an explicit and delayed acknowledgement to an execution-requesting client if a workflow is completed by a first workflow engine. In Nakamura, if a document is aborted, the corresponding participant is notified (Para. 0097); however, Appellant could find no disclosure in Nakamura regarding sending an acknowledgement to the client if a workflow is completed.

Further, as mentioned above in the summary of Nakamura, Nakamura discloses only a single workflow engine (server 110), so Nakamura cannot assign the workflow to a second workflow engine.

No proper combination of Campbell and Nakamura can provide these missing elements.

Therefore, a *prima facie* case of obviousness has not been established by the combination of Campbell and Nakamura. Therefore, it is respectfully asserted that no proper combination of Campbell and Nakamura can be made to disclose the embodiments claimed, and it is respectfully requested that the rejection of claims 1-36 under 35 U.S.C. § 103(a) be reconsidered and withdrawn.

8. SUMMARY

It is respectfully submitted that no *prima facie* case of obviousness under 35 U.S.C. §103 has been established by the Office. Therefore, it is respectfully requested that the rejections of claims 1-36 be reconsidered and withdrawn. The Appellant respectfully submits that all of the claims are in condition for allowance, and notification to that effect is earnestly requested. The Examiner is invited to telephone the Appellant's attorney, Walter W. Nielsen at (602) 298-8920, or the undersigned attorney at (612) 349-9592, to facilitate prosecution of this Application. If necessary, please charge any additional fees or credit overpayment to Deposit Account No. 19-0743.

Respectfully submitted,

MINGQIU SUN ET AL.

By their Representatives,

SCHWEGMAN, LUNDBERG, WOESSNER &
KLUTH, P.A.

Attorneys for Intel Corporation

P.O. Box 2938

Minneapolis, Minnesota 55402

(612) 349-9592

Date Oct. 11, 2005

By Ann M. McCrackin
Ann M. McCrackin
Reg. No. 42,858

CERTIFICATE UNDER 37 CFR 1.8: The undersigned hereby certifies that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail, in an envelope addressed to: Mail Stop Appeal Brief - Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on this 11th day of October, 2005.

Carolyn Hulsey
Name

Carolyn Hulsey
Signature

CLAIMS APPENDIX

1. (Rejected) A method to be performed by a data processing system comprising:
 providing distributed queuing of workflows, whose execution is requested by one
 or more execution-requesting clients, among a plurality of workflow engines; and
 if a workflow is completed by a first workflow engine for an execution-requesting
 client, sending an explicit and delayed acknowledgement to the execution-requesting
 client, else assigning the workflow to a second workflow engine.

2. (Rejected) The method recited in claim 1, wherein providing is performed by a
 load manager.

3. (Rejected) The method recited in claim 2, wherein the load manager comprises a
 commercially available middleware product.

4. (Rejected) The method recited in claim 1, wherein the explicit and delayed
 acknowledgement is performed by a certified messaging capability.

5. (Rejected) The method recited in claim 4, wherein the certified messaging
 capability is performed by a load manager.

6. (Rejected) The method recited in claim 4, wherein the load manager comprises a
 commercially available middleware product.

7. (Rejected) The method recited in claim 4, wherein the certified messaging
 capability is performed by a certified message receiver forming part of the workflow.

8. (Rejected) The method recited in claim 4 and further comprising:
the certified messaging capability sending an explicit and delayed acknowledgement to the execution-requesting client if the workflow is completed by the second workflow engine.
9. (Rejected) A method to be performed by a computer network comprising a plurality of clients and a plurality of workflow engines:
providing distributed queuing of workflows, whose execution can be requested by one or more execution-requesting clients, among the plurality of workflow engines; and
determining whether a workflow has been completed by a first workflow engine on behalf of an execution-requesting client; and
if so, sending an explicit and delayed acknowledgement to the execution-requesting client;
otherwise, assigning the workflow to a second workflow engine.
10. (Rejected) The method recited in claim 9, wherein providing is performed by a load manager.
11. (Rejected) The method recited in claim 10, wherein the load manager comprises a commercially available middleware product.
12. (Rejected) The method recited in claim 9, wherein sending is performed by a certified messaging capability.
13. (Rejected) The method recited in claim 12, wherein the certified messaging capability is performed by a load manager.
14. (Rejected) The method recited in claim 12, wherein the load manager comprises a commercially available middleware product.

15. (Rejected) The method recited in claim 12, wherein the certified messaging capability is performed by a certified message receiver in the workflow.

16. (Rejected) The method recited in claim 12 and further comprising:
the certified messaging capability sending an explicit and delayed acknowledgement to the execution-requesting client if the workflow is completed by the second workflow engine.

17. (Rejected) A computer adapted for use in a computer network comprising a plurality of workflow engines, the computer executing a computer program, the computer program operating the computer in a fault-tolerant manner and comprising the operations of:

requesting a workflow execution on behalf of a client;
a distributed queuing capability assigning the workflow execution to a first workflow engine;
determining whether the workflow execution has been completed by the first workflow engine; and
if so, sending an explicit and delayed acknowledgement to the client;
otherwise, assigning the workflow execution to a second workflow engine.

18. (Rejected) The computer recited in claim 17, wherein requesting is performed by a load manager.

19. (Rejected) The computer recited in claim 17, wherein sending is performed by a certified messaging capability.

20. (Rejected) The computer recited in claim 19, wherein the certified messaging capability is performed by a certified message receiver in the first workflow engine.

21. (Rejected) The computer recited in claim 19 and further comprising:
the certified messaging capability sending an explicit and delayed
acknowledgement to the client if the workflow execution is completed by the second
workflow engine.
22. (Rejected) A computer network comprising:
a plurality of clients;
a plurality of workflow engines; and
at least one computer program, the computer program operating in a fault-tolerant
manner and performing the operations of:
requesting a workflow execution on behalf of a client;
assigning the workflow execution to a first workflow engine;
determining whether the workflow execution has been completed by the first
workflow engine; and
if so, sending an explicit and delayed acknowledgement to the client;
otherwise, assigning the workflow execution to a second workflow engine.
23. (Rejected) The computer network recited in claim 22, wherein requesting is
performed by a load manager having a distributed queuing capability.
24. (Rejected) The computer network recited in claim 22, wherein sending is
performed by a certified messaging capability.
25. (Rejected) The computer network recited in claim 24, wherein the certified
messaging capability is performed by a certified message receiver in the first workflow
engine.

26. (Rejected) The computer network recited in claim 24 and further comprising:
the certified messaging capability sending an explicit and delayed
acknowledgement to the client if the workflow execution is completed by the second
workflow engine.
27. (Rejected) A computer-readable medium containing computer instructions for
instructing a processor, the processor adapted for use in a computer network comprising a
plurality of workflow engines, wherein the instructions comprise:
requesting a workflow execution on behalf of a client;
a distributed queuing capability assigning the workflow execution to a first
workflow engine;
determining whether the workflow execution has been completed by the first
workflow engine; and
if so, sending an explicit and delayed acknowledgement to the client;
otherwise, assigning the workflow execution to a second workflow engine.
28. (Rejected) The computer-readable medium recited in claim 27, wherein
requesting is performed by a load manager.
29. (Rejected) The computer-readable medium recited in claim 27, wherein sending
is performed by a certified messaging capability.
30. (Rejected) The computer-readable medium recited in claim 29, wherein the
certified messaging capability is performed by a certified message receiver in the first
workflow engine.

31. (Rejected) The computer-readable medium recited in claim 29 and further comprising:

the certified messaging capability sending an explicit and delayed acknowledgement to the client if the workflow execution is completed by the second workflow engine.

32. (Rejected) An article comprising a machine-accessible medium having instructions for instructing a processor forming part of a plurality of workflow engines, wherein the instructions, when accessed, result in a machine performing:

requesting a workflow execution on behalf of a client;

assigning the workflow execution to a first workflow engine;

determining whether the workflow execution has been completed by the first workflow engine; and

if so, sending an explicit and delayed acknowledgement to the client;

otherwise, assigning the workflow execution to a second workflow engine.

33. (Rejected) The article recited in claim 32, wherein requesting is performed by a load manager having a distributed queuing capability.

34. (Rejected) The article recited in claim 32, wherein sending is performed by a certified messaging capability.

35. (Rejected) The article recited in claim 34, wherein the certified messaging capability is performed by a certified message receiver in the first workflow engine.

36. (Rejected) The article recited in claim 34 and further comprising:
the certified messaging capability sending an explicit and delayed acknowledgement to the client if the workflow execution is completed by the second workflow engine.

EVIDENCE APPENDIX

NONE.

RELATED PROCEEDINGS APPENDIX

No Related Proceedings are known to the Appellants' Representative.